On the roles of markedness and contradiction in the use of alternatives*

Roni Katzir
rkatzir@post.tau.ac.il

September 5, 2013

Scalar implicature (SI) is often viewed as reasoning from what was said to what was left unsaid. For example, a speaker who utters *John has three children* can license the SI that John has exactly three children, not four or more. This is often attributed to the fact that, if the speaker had thought that John had more than three children, the speaker could have conveyed this information by saying *John has four children*. The speaker didn't, so we conclude that John doesn't. Taking into account alternative statements that were not used is an attractive idea, but it requires addressing several non-trivial challenges, such as where in the cognitive architecture the relevant reasoning about such unsaid alternatives is performed and how exactly the alternatives are handled.

An elementary but still important challenge is the characterization of what counts as an alternative: depending on how the reasoning process is implemented, different choices of alternatives can lead to a given inference, to its negation, to no inference, or to contradiction. Since SIs are by and large an orderly, systematic family of inferences, a proper theory of SIs will have to provide a general characterization of the alternatives that correctly predicts the inferences that are actually observed.

SI is not alone in making reference to alternatives. Indeed, alternative-sensitive computations have been central to accounts of morpho-phonology (Anderson, 1969; Kiparsky, 1973; Halle and Marantz, 1993; Halle, 1997), the semantics of focus and questions (Jackendoff, 1972; Hamblin, 1973; Rooth, 1985), and recently even syntax (Chomsky, 1995; Fox, 2000). Many of the questions that arise with SI have counterparts in these other domains. Here I will discuss SI alongside two alternative-sensitive phenomena that are particularly close to it: the semantics of association with focus (AF) in the case of the exhaustive operator *only*; and the appropriateness conditions on what is sometimes referred to as free focus (FF).

Despite suggestive similarities between SI, AF, and FF, the literature has, for the most part, treated them as three different phenomena. In particular, proposals for the alternatives for each process have often been quite different from one another. Here I will

*Acknowledgments: This paper builds on several previous papers, and I wish to thank all those who are acknowledged in those works. Special thanks for their help with the current paper to Luka Crnič, Danny Fox, Tova Friedman, Yael Greenberg, Yael Sharvit, and Raj Singh. This work has been supported by ISF grant 187/11.

defend the position that the three phenomena are in fact very closely related. Specifically, I will argue, building on Katzir (2007), Fox and Katzir (2011), and Katzir (2013), that the alternatives for the three processes are identical and involve markedness-based substitutions within focused constituents, where markedness is cashed out in terms of structural complexity.

1 Alternatives for implicature and focus

1.1 Background: SI, AF, FF

Consider the following simple sentences with their SIs:[1]

(1) Scalar Implicature (SI):

a. John did *some* of the homework
SI: ¬ [John did *all* of the homework]

b. John did the reading *or* the homework
SI: ¬ [John did the reading *and* the homework]

In (1a), the assertion that John did some of the homework is strengthened by conjoining it with the SI that it is not the case that John did all of the homework. This SI is derived by negating *John did all of the homework*, an alternative of the assertion in which *some* is replaced with *all*. In (1b), the assertion that John did the reading or the homework is strengthened by conjoining it with the SI that it is not the case that John did the reading and the homework. This SI is derived by negating *John did the reading and the homework*, an alternative of the assertion in which *or* is replaced with *and*. More generally, SIs are obtained by negating alternatives, sometimes called *scalar alternatives*, that are variants of the assertion in which certain elements have been replaced with others.

AF, as in (2), is similarly dependent on alternatives:[2]

(2) Association with Focus (AF):

a. John only introduced Mary$_F$ to Sue
$\Rightarrow$ ¬ John introduced Jane to Sue

b. John only introduced Mary to Sue$_F$
$\Rightarrow$ ¬ John introduced Mary to Jane

A sentence of the form *only*(S) has entailments that are similar to the SIs of S when uttered in isolation. In (2a), *only* strengthens the prejacent by conjoining it with the proposition that it is not the case that John introduced Jane to Sue.[3] This proposition

[1] The brackets in examples (1a) and (1b) are included to mark the scope of negation. To enhance readability, I will omit the brackets in subsequent examples, but the scope of negation will still be maximal.

[2] Here and in subsequent examples I will use the subscript F to mark the placement of focus. When the relevant focus is placed on a constituent that spans more than one word, its scope will be marked with brackets.

[3] AF is treated here as applying at the propositional level, rather than at the level of properties as in Rooth's account. This is done for expository convenience only (specifically, in order to highlight the similarity with SI). I also gloss over important questions of the division of labor between assertion and presupposition in the semantics of *only*.

is the negation of *John introduced Jane to Sue*, an alternative of the prejacent in which the focus-marked *Mary* is replaced with *Sue*. Alternatives that are derived through substitutions within focused constituents are sometimes called *focus alternatives*. In (2b), *only* strengthens the prejacent by conjoining it with the proposition that it is not the case that John introduced Mary to Jane. This proposition is the negation of *John introduced Mary to Jane*, an alternative of the prejacent in which the focus-marked *Sue* is replaced with *Jane*. More generally, *only* negates alternatives that are variants of the prejacent in which certain elements have been replaced with others.

The conditions in which focus is appropriate in context – the so-called *free* use of focus (FF) – are also sensitive to alternatives. For example, in the context of (3), it is appropriate to respond with (3a), in which the subject *John* is focused, but it is inappropriate to continue with (3b), in which it is the verb *drinks* that is focused.

(3) Mary drinks coffee
  a. And John$_F$ drinks coffee
  b. # And John drinks$_F$ coffee

The pattern in (3) is usually derived by requiring that the context contain – or support the accommodation of – a focus-alternative of the utterance:

(4) APPROPRIATENESS CONDITION ON FREE FOCUS: each sentence must have a focus-alternative in the context

Since *John* is focused in (3a), its focus alternatives are derived by replacing it with other elements, such as *Mary*, so (3) is a focus alternative. In (3b), on the other hand, it is the verb that is focused, and so its focus alternatives are of the form *John X coffee*, and nothing in the context is of this form.[4]

1.2 Formal alternatives and context

Call the set of alternatives that end up being used A. Since Horn (1972), A has been taken to be determined by an interaction of both formal and contextual factors. Formal factors map a sentence S to a set of formal alternatives, call it $F(S)$. Contextual factors, which we can think of as a second set of alternatives, C, then help in narrowing $F(S)$ down to A. In section 5 we will try to obtain a clearer picture of the interaction of F and C. For the moment, let us assume that the actual alternatives are simply the intersection of the two:

(5) $A = C \cap F(S)$

[4]There are additional factors that govern the placement of focus. For example, *And John$_F$ drinks$_F$ coffee* satisfies (4) in the context of (3): it has focus on the subject *John* (along with the verb *drinks*) and consequently has the context-setting sentence *Mary drinks coffee* in (3) as an alternative; nevertheless, it is an inappropriate response to (3). Following Schwarzschild (1999), this is usually accounted for in terms of economy: focusing the subject alone, as in (3a), is sufficient, and hence any additional focus marking is ruled out. See Schwarzschild (1999), Sauerland (2005), Wagner (2005), Spathas (2010), and Büring (2012) for additional constraints of this kind and much relevant discussion. In section 6 we will see an argument for a different factor in the acceptability of FF, namely its ability to contribute to a (sometimes exhaustified) meaning that addresses a previous expectation.

C is invoked since, as already observed by Horn (1972, p. 112) for SIs and by Rooth (1985, pp. 42–3) for AF, different inferences are generated in different contexts. For example, (6), repeated from (1a) above, has (6a) and (6b) as optional SIs. And the sentence in (7), uttered in response to the question *What did Kim do yesterday?*, can be taken to entail either (7a) or (7b), among other possibilities, depending on context.

(6) John did *some* of the homework
   a. SI? ¬ John did *most* of the homework
   b. SI? ¬ John did *much* of the homework

(7) Kim only [wrote a poem]$_F$
   a. ⇒? ¬ Kim saw a movie
   b. ⇒? ¬ Kim ate an apple

There are two main reasons for restricting the alternatives using the formally defined $F(S)$. One is the so-called symmetry problem, which I will discuss in the following sections. A second, more obvious reason is that we need $F(S)$ in order to account for focus sensitivity. For example, (2a) and (2b) above are identical other than the placement of focus, and so their different entailments cannot be attributed to context alone. The FF examples in (3) make a similar point.

(8) $F(S) = \{S' : S'$ *is derived by replacing focused items in* S *with allowable substitutions*$\}$

The focus-sensitivity of the formal alternatives is clearest in AF and FF, but as noted by Rooth (1992), SI is focus sensitive as well:

(9) How did the exam go?
   a. Well, I passed$_F$
   b. Well, I$_F$ passed

In (9a) focus falls on the VP *passed*, and this response licenses the inference, presumably an SI, that I did not do any better than passing. No inference about who else may have passed is licensed in this case. In (9b), on the other hand, focus falls on the subject *I*, and this response licenses the inference that someone else did not pass. No inference about whether or not I did better than passing is licensed. The restriction of F in (8) according to focus marking is thus required not just for AF and FF but also for SI.

1.3 Outline

While many accounts of SI, AF, and FF can agree on the general schema in (8), there are considerable differences in what these accounts consider the allowable substitutions to be. In what follows I will argue for a unified account of all three phenomena. Building on Katzir (2007), Fox and Katzir (2011), and Katzir (2013), I will argue that the allowable substitutions for all three phenomena are defined in terms of structural

complexity. Before investigating the details of the complexity approach, it will be convenient to examine a more straightforward approach, familiar from the literature on focus semantics, in which substitutions are based on semantic type. As we will see in section 2, the definition in terms of semantic type admits too many alternatives, leaving it vulnerable to the symmetry problem mentioned above. In section 3 I will present the complexity approach, a general definition that is capable of breaking symmetry in the correct way.

In section 4 I will compare the complexity approach with a definition of substitutions in terms of the semantic condition of monotonicity. Monotonicity is a restrictive approach: in order to account for the systematic absence of certain inferences, it is sparing in the alternatives that it admits. We will see, however, that there is a different way to account for the missing inferences, one that makes it possible to maintain the complexity approach: we could admit the additional alternatives that the complexity approach allows and that monotonicity rules out, and we could attribute the lack of the relevant inferences to a contradiction that arises due to these additional alternatives. In order to choose between complexity, which accounts for the pattern in terms of too many alternatives, and monotonicity, which accounts for the pattern in terms of too few alternatives, we need to find environments in which the additional alternatives brought in under the complexity approach do not give rise to contradiction. As we will see, when we do this the missing alternatives will emerge, thus supporting the complexity account.

An interesting consequence of our findings from the comparison of the structural approach with monotonicity is that, both in SI and in AF, context is quite limited in its ability to select alternatives. We will examine a dilemma for alternative-selection schemes. This dilemma, a generalization of the notion of symmetry, involves cases in which negating all the alternatives would be inconsistent with the assertion but choosing from among the alternatives would be arbitrary. We will see that when this dilemma arises, context is incapable of helping. I will discuss this consequence in section 5.

Combined with section 3, which tries to show that F is the same for both SI and AF, section 5, which tries to show that contextual restriction is the same for both processes, will complete the present attempt to argue for a unification of the two phenomena. This still leaves FF as a potentially unrelated alternative-sensitive phenomenon. In section 6, however, I will try to show that this is not the case, and that FF is closely related to SI and AF. I will start by considering yet another restrictive semantic characterization of the allowable substitutions, this time in terms of contrast. As with monotonicity, the motivation for contrast will come from inferences that are systematically missing; and as in the discussion of monotonicity, I will argue that contrast is irrelevant, and that the missing inferences should be accounted for in terms of too many alternatives and not in terms of too few alternatives. Significantly, the pattern that motivated the contrast proposal involves both AF and FF. And while accounting for the pattern in the case of AF will be a very close repetition of the argument against monotonicity in section 4, extending the account to FF will require replacing the appeal to the appropriateness condition on FF in 4 with an appeal to an exhaustification operation, along the lines of SI and AF. For the reasons discussed earlier in the paper, such a process suggests the same structurally-defined alternatives as for SI and AF, thus completing the argument for the unified treatment of all three processes.

2 Semantic type and the symmetry problem

A common view in the domain of focus semantics and pragmatics (see Jackendoff 1972, pp. 240–247 and Rooth 1985, pp. 13–14) takes the allowable substitutions in (8) to be characterized by semantic type:

(10) *Allowable substitutions* (TYPES): elements of the same semantic type

In (2a) above (= *John only introduced Mary$_F$ to Sue*), for example, the single focus in the sentence is on *Mary*, whose semantic type is e. Consequently, $F(2a)$ will be the set of variants of the prejacent of (2a) in which *Mary* has been replaced with some other individual: $F(2a)$ = {*John introduced Mary to Sue*, *John introduced Sue to Sue*, *John introduced Jane to Sue*, ...}. In (7) (= *Kim only [wrote a poem]$_F$*), the focus falls on the entire *VP*, which is of type $< e, t >$. Consequently, $F(7)$ will be the set of variants of the prejacent of (7) in which *wrote a poem* has been replaced with some other property: $F(7)$ = {*Kim wrote a poem*, *Kim ate an apple*, *Kim slept*, ...}.

The characterization of allowable substitutions according to TYPES, as defined in (10), is general and permissive. In fact, it is too permissive: the imposition of no further restriction beyond semantic type exposes TYPES to an insidious problem, due to Kroch (1972) and developed further by von Fintel and Heim (1997) and Horn (2000), and often referred to as the *symmetry problem*. To facilitate our discussion, let us define symmetry as follows:[5]

(11) SYMMETRY (NARROW SENSE; TO BE GENERALIZED IN (38)): Sentences S_1 and S_2 are *symmetric* alternatives of S if both
   a. $[\![S_1]\!] \cup [\![S_2]\!] = [\![S]\!]$, and
   b. $[\![S_1]\!] \cap [\![S_2]\!] = \emptyset$

For (6) (= *John did some of the homework*), for example, the following are symmetric alternatives:

(12) Potential alternatives for (6):
   a. John did *all* of the homework
   b. John did *some but not all* of the homework

More generally, given any S and a stronger S_1, S_1 and $S_2 = S \wedge \neg S_1$ will be symmetric alternatives. The SI in (6) was generated by taking only (12a) into account. If we ignored it and took only (12b) into account we would infer the opposite, namely that John did all of the homework, which is of course the wrong result. If we took both alternatives into account, we would not be able to negate them without arriving at a contradiction. The challenge is to ensure that of the two symmetric potential alternatives, only S_1 ends up being negated, so as to derive the observed SI. More generally, the symmetry problem is the challenge of ensuring that symmetry is broken in a way that matches the observed pattern of inference.

[5]This is a somewhat narrower definition than is needed, but for the present it will do. In section 5 we will see a more general notion of symmetry.

As noted in Fox and Katzir (2011), the symmetry problem arises not just for SI but also for AF.[6] Consider again (7) above (= *Kim only [wrote a poem]$_F$*), along with its two symmetric type-based alternatives in (13).

(13) Potential alternatives for (7):
   a. S_1: Kim wrote a poem and saw a movie
   b. S_2: Kim wrote a poem and didn't see a movie

Both in SI and in AF, then, we are faced with a similar challenge. In order to derive the correct pattern of inference, we must be able to break symmetry by negating S_1 but not S_2. In the literature on SI, this symmetry breaking is often obtained within F by defining the allowable substitutions in terms of *Horn Scales*, a family of lexical stipulations. Scales, which were introduced by Horn (1972) and further studied by Gazdar (1979), Atlas and Levinson (1981), Hirschberg (1985/1991), and others, allow us to state, for example, that *some* can be replaced by *all* but not by *some but not all*, thus deriving the correct result for (6) and its potential alternatives in (12). While successful in this case, lexical stipulations do not provide the basis for a predictive theory that can handle arbitrary constituents, as in (13). What we need is a general characterization of allowable substitutions.

The definition of Types in (10) does provide a general characterization of allowable substitutions, but it is a characterization that ensures that both S_1 and S_2 will be in F, leaving symmetry breaking for C. This, however, is problematic, since it predicts, incorrectly, that we should be able to reverse the standard pattern of entailment by making S_2 but not S_1 salient:

(14) Mary read exactly three books. What did John do?
   a. John read three books
      *SI ¬ John read exactly three books
   b. John only [read three books]$_F$
      ⇏ ¬ John read exactly three books

In (14), context makes *read exactly three books* salient. If C could break symmetry, we would expect it to be able to keep *read exactly three books* and not (the unmentioned) *read four books*. This would derive the inference (an SI in (14a) and an entailment in (14b)) that John read more than three books, an impossible inference.[7] It seems, then, that F in both (14a) and (14b) must be asymmetric to begin with, which in turn means that Types is incorrect.

3 Complexity

Can we let F break symmetry while maintaining a general definition? One approach, with a history going back to Zipf (1949), Grice (1989), McCawley (1978), and Horn

[6] The relevance of symmetry for FF is less obvious, but in section 6 I will try to show that symmetry is central for FF as well.

[7] And if that inference were available, it would be hard to see why (14a) and (14b) should not be perfectly acceptable answers to the question in (14). In fact, they are decidedly odd in that role, thus further highlighting the puzzle for Types.

(1984), allows us to break symmetry based on considerations of complexity: *some* and *all* are of (roughly) the same length; *some but not all* is longer.[8] This suggests the following:

(15) *Allowable substitutions* (COMPLEXITY; SCHEMATIC VERSION; TO BE REVISED IN (25)): elements of (up to) the same complexity

I will try to defend a version of COMPLEXITY, but first we must make the statement in (15) more precise. In particular, we should decide whether to use *up to* or not, and we should also have a clearer understanding of how complexity is defined.

Regarding *up to*, I will follow Horn (1984, 2000) in concluding that strictly simpler alternatives are generally available. A nagging worry, however, is the scarcity of implicatures that are based on negating strictly simpler alternatives. This scarcity, however, becomes less surprising when one notes that complexity often comes from modifiers, which also strengthen the meaning of the sentence. Often, then, a simpler alternative will be a weaker alternative and thus not negatable. The following examples illustrate this point.

(16) A tall candidate was interviewed ($= \phi$)
*SI: $\neg$ A candidate was interviewed ($= \phi'$)

(17) I only interviewed [a tall candidate]$_F$ ($= \psi$)
$\not\Rightarrow \neg$ I interviewed a candidate ($= \psi'$)

Since $[\![\phi]\!] \subset [\![\phi']\!]$, we cannot negate ϕ' without contradicting the assertion ϕ. Similarly for ψ and ψ'. In order to test for inferences that are based on simpler alternatives, we must look at places where simpler is not weaker. One straightforward way to do so is to consider downward-entailing environments, which reverse entailment relations. Here are a few examples:[9]

(18) a. Every tall candidate got interviewed
SI: $\neg$ Every candidate got interviewed
b. John doubts that many dogs with long tails will be sold
SI: $\neg$ John doubts that many dogs will be sold
c. John is forbidden to sing loudly
SI: $\neg$ John is forbidden to sing

(19) a. I only interviewed [every tall candidate]$_F$
$\Rightarrow \neg$ I interviewed every candidate
b. John only doubts that many dogs with long tails will be sold
$\Rightarrow \neg$ John doubts that many dogs will be sold

[8] The absence of single-word versions of *not all* and *some but not all* is part of a systematic pattern across languages and categories that was discovered and analyzed by Horn (1972). See Katzir and Singh (2013) for further discussion.

[9] Another place to look is disjunctions: ϕ and ψ will always be at least as strong as – and in most felicitous cases strictly stronger than – the disjunction $\phi \vee \psi$. If strictly simpler alternatives are available, we would expect a disjunction $\phi \vee \psi$ to show the effects of having ϕ and ψ as alternatives. See Sauerland (2004b), Sauerland (2004a), and Fox (2007) for evidence that this is indeed the case.

c. John is only forbidden to sing loudly
   $\Rightarrow \neg$ John is forbidden to sing

It seems, then, that the *up to* part in the schematic definition of COMPLEXITY in (15) is needed: strictly simpler alternatives are possible.

A slight variation on the latest examples shows that, while *tall candidate* has *candidate* as an alternative, the opposite is not true:

(20) Every boy spoke to a candidate
   *SI: $\neg$ Every boy spoke to a tall candidate

(21) I only introduced every boy to a candidate
   $\nRightarrow \neg$ I introduced every boy to a tall candidate

In other words, the relation `alternative-of` is not symmetric: there are cases in which ϕ is an alternative of ψ but ψ is not an alternative of ϕ.[10] This conclusion, which is very much in line with Horn (1984) and Horn (2000), is further evidence against TYPES, defined in (10) above, but is relevant for additional theories of allowable substitutions, including the monotonicity-based approach that we will examine in section 4 below.

As for making complexity more precise, many intuitive notions of complexity, such as phonetic effort, syllable count, etc., have long been known not to work. For example, the monosyllabic *some* is generally taken to have the bisyllabic *every* as an alternative.[11]

A notion of complexity that initially seems to avoid the pitfalls for markedness-based approaches is a syntactic one. We can consider X' simpler than X if it can be derived from X by successive replacements of sub-constituents of X with elements of the lexicon or their own sub-constituents. For the DP *every tall candidate* in (18a) and (19a), for example, the simpler DP *every candidate* can be obtained by substituting *candidate* for *tall candidate* (reassuringly, there is no similar way to get from *every candidate* to *every tall candidate*). Syntactic complexity, defined through substitutions, allows us to break symmetry in all the cases examined so far. Importantly, it also avoids the challenges to measures of phonetic effort and syllable count: since *every* is in the lexicon, it can be substituted for *some* regardless of the additional syllable obtained in the process.

Matsumoto (1995) provides an example that would seem to go against the very notion of complexity:

(22) It was warm yesterday, and it is a little bit more than warm today
   SI: $\neg$ It was a little bit more than warm yesterday

[10] The notion of symmetry just mentioned is the familiar set-theoretic notion of symmetry for relations. This notion, which will only be relevant very locally here and in the beginning of section 4, is unrelated to the alternative-specific notion of symmetry in (11), a notion that will be relevant throughout the paper and will be generalized in (38). Despite the unfortunate terminological overlap, I hope context will prevent confusion from arising.

[11] See also related discussions in the literature on lexical and morphological blocking, especially Horn (1978, 1984) and Poser (1992).

The SI in (22) relies on substituting the long *a little bit more than warm* for the shorter *warm*, but this seems to indicate that longer alternatives are available.

We can address Matsumoto's challenge by relying on the fact that the longer alternative was made salient in the context of (22). If we extend the allowable substitutions – earlier specified as coming from the lexicon and from sub-constituents – so as to also include contextually salient constituents, we will have a notion of complexity that can account for (22). In total:

(23) $X' \precsim_C X$ if X' can be derived from X by successive steps in which a sub-constituent z of X is replaced with an element of the substitution source for z in C, $SS(z, C)$

(24) $SS(z, C)$, the substitution source for z in context C, is the union of:

a. The lexicon

b. The sub-constituents of z

c. The set of salient constituents in C (that is, constituents of the structures of utterances made in recent discourse)

(25) *Allowable substitutions* (COMPLEXITY; FINAL VERSION): Allowable substitutions for X in context $C = \{X' : X' \precsim_C X\}$

With (25) we have arrived at a usable complexity-based characterization of allowable substitutions. We have just seen evidence that COMPLEXITY, as defined in (25), is the correct basis for the definition of F both for SI and for AF (evidence that F is defined in the same way also for FF will be less direct and will wait until section 6). In the following section we will see further evidence supporting COMPLEXITY, and in section 5 we will see evidence that SI and AF share not only the same definition of F but also the same form of contextual restriction.

4 Monotonicity

Structure is not the only systematic way to derive symmetry breaking in F. A prominent semantic approach, based on Horn (1989) and Matsumoto (1995), uses the following observation: *some* and *all* are both upward monotone in their second argument; *some but not all*, on the other hand, is non-monotone. This suggests the following characterization of allowable substitutions:[12]

(26) *Allowable substitutions* (MONOTONICITY; PRELIMINARY VERSION; TO BE REVISED IN (27)): elements of the same semantic type and same monotonicity

As we already saw from the contrast between (18) and (19) on the one hand and (20) and (21) on the other hand, the `alternative-of` relation is not symmetric, so (26) on its own does not suffice. However, it is still possible to maintain that monotonicity is operative alongside complexity. This, I take it, is the view advocated by Swanson (2010).

[12]Neither Horn (1989) nor Matsumoto (1995) is committed to this being a sufficient condition, only a necessary one.

(27) *Allowable substitutions* (MONOTONICITY; FINAL VERSION; INCLUDES COMPLEXITY): (26) ∩ (25)

COMPLEXITY (without monotonicity; as in (25)) and MONOTONICITY (that is complexity with monotonicity, as in (27)) both account for all of the data we have seen so far. They do, however, make divergent predictions. According to COMPLEXITY, non-monotone elements can have alternatives and be themselves alternatives of other elements. According to MONOTONICITY, non-monotone elements cannot stand in the alternative-of relation with monotone elements. An observation due to Danny Fox and reported in Katzir (2007) seems to support MONOTONICITY: when Matsumoto's example is modified to involve non-monotone elements instead of the original monotone ones, no inference arises:

(28) John did *some* of the homework yesterday, and he did *just some* of the homework today
*SI: ¬ John did *just some* of the homework yesterday

According to our definition of structural complexity in (23), the presence of *just some* in (28) makes *just some* an allowable substitution for *some* in that context, which in turn makes the absence of the relevant SI puzzling. For MONOTONICITY, on the other hand, this absence is predictable: *just some* is non-monotone and is therefore an impossible replacement for the (upward-monotone) *must*, regardless of complexity. Taken together, this would seem like an argument for MONOTONICITY and against COMPLEXITY on its own. However, as further discussed in Katzir (2007), there is a confound in this argument, due to symmetry, which prevents us from using (28) to choose between COMPLEXITY and MONOTONICITY. In our definition of complexity, we had to allow substitutions from the lexicon. This means that COMPLEXITY predicts (at least) two alternatives for *John did some of the homework yesterday* in the context of (28), one derived by replacing *some* with *just some* and another derived by replacing *some* with *all*:

(29) a. John did *just some* of the homework yesterday ...
b. John did *all* of the homework yesterday ...

(29a) and (29b) are symmetric alternatives of (28), and COMPLEXITY predicts that both are in $F(28)$. If symmetry in F poses a problem for SI – we will shortly make this more precise – then we can hope to account for the absence of the relevant SI in (28) without relying on monotonicity: the alternative (29a) is available, but so is (29b), and since we cannot negate both without contradicting the assertion, neither will be negated, and the unavailability of the SI $\neg(29a)$ is accounted for. That is, while MONOTONICITY accounts for the missing SI in terms of too few alternatives, COMPLEXITY accounts for it in terms of too many alternatives.

In order to tease apart the predictions of MONOTONICITY and COMPLEXITY, we can construct a variant of (28) in which the counterparts of the alternatives in (29) can be negated without contradicting the assertion. In such a configuration, MONOTONICITY predicts that the alternative with *just some* will still be unavailable (since it is still non-monotonic), while COMPLEXITY predicts that both alternatives will now be negatable. Embedding under a universal operator allows us to create exactly this kind of test (cf. Sauerland, 2004a; Fox and Hackl, 2006; Fox, 2007):

(30) John was required to do *some* of the homework yesterday, and he was required to do *just some* of the homework today
SI: ¬ John was required to do *just some* of the homework yesterday
SI: ¬ John was required to do *all* of the homework yesterday

In (30) we modified (28) to involve embedding under the universal operator *require*, and we can observe that the SI with the non-monotone *just some* emerges. We can conclude that monotonicity plays no role in SIs; the problem is one of too many alternatives. The exact same reasoning can be repeated with AF.[13]

5 Where symmetry can be broken

As noted by Fox and Katzir (2011), the paradigm we just saw as part of the argument against monotonicity allows us to sharpen a generalization that we discussed earlier, in section 2. There, we saw that symmetry, in the sense of definition (11), must sometimes be broken by F, an observation that went against inherently symmetric theories of F, such as Rooth (1985)'s. We can now see that it is not just that F is one place where symmetry can be broken: F is the *only* place where symmetry is ever broken. If F happens to contain a symmetric subset X, no element of X will be negated. In particular, we can conclude that context is incapable of breaking symmetry.

The inability of context to break symmetry is somewhat surprising. It is certainly unexpected if contextual restriction is taken to be a selection of salient alternatives: there is nothing to prevent discourse from making exactly one of two symmetric alternatives salient; in fact, we already saw several examples in which discourse does just that. But in all these cases, the salience of just one of the symmetric alternatives was not able to break symmetry. The conclusion from this is not that context does not affect the choice of alternatives – we saw, through examples (6) and (7) above, that context plays an important role in this choice – but rather that contextual restriction is limited in how it affects the selection of alternatives, and that this limitation prevents it from breaking symmetry. Given this conclusion, it might be useful to take a closer look at the notion of symmetry, which is what I will try to do in the present section.

The definition of symmetry in (11) captures a clear and particularly useful case of a potentially lethal dilemma for alternative-selection schemes: S_1 and S_2 cannot both be negated consistently with S, and any choice between the two alternatives would be arbitrary. Where we find inferences based systematically on one of the two alternatives, we must rely on a formal criterion – we considered complexity and monotonicity and saw evidence for the former – to obtain this result. Dilemmas of this kind are a more general phenomenon than (11) suggests, however, and in this section I will try to show that the pattern we saw above can be repeated with the more general notion. This generalization, which will lead us to Fox (2007)'s notion of innocent exclusion, was not needed for the arguments we saw above in favor of Complexity, as defined in (25),

[13] Of course, it is quite possible that other factors do play a role. For example, Matsumoto (1995), Geurts (2010), and Swanson (2010) discuss cases in which certain lexical substitutions are unavailable, possibly due to what is considered acceptable level of detail. Complexity is entirely compatible with additional requirements. The argument above only shows that monotonicity is not one of those requirements.

and against both TYPES, as defined in (10), and MONOTONICITY, as defined in (27). It does have some interesting architectural consequences, though, which we will see shortly. It will also be relevant for an argument against a role for contrast, which we will see in section 6.

There are three independent senses in which (11) offers too narrow a characterization of the dilemma for alternative-selection schemes. I will start by briefly presenting each of the three ways in which we may wish to generalize (11). We will then see that when F contains alternatives that satisfy symmetry in the generalized sense, none of them is negated. This will generalize what we saw in the argument against monotonicity and indicate that context is incapable of breaking not just symmetry as defined in (11) but generalized symmetry more broadly.

A first sense in which (11) is too limited is that it considers only two alternatives. For example, *John died last weekend* has two symmetric potential alternatives, *John died last Saturday* and *John died last Sunday*, and there is no way to choose between them in a non-arbitrary way. But refining the partition of S into additional mutually exclusive subsets will do little to help us select an alternative in a non-arbitrary way. For example, *John died last week* has seven potential alternatives, one for each day of the week, and it doesn't seem to make the choice any less arbitrary than in the earlier case with just two alternatives. We may wish to extend (11) so as to capture cases in which the set of problematic alternatives has more than two members.

A second sense in which (11) is too limited is its requirement that S_1 and S_2 be disjoint. *John died last weekend* has the two symmetric potential alternatives just mentioned, *John died last Saturday* and *John died last Sunday*. But now consider the minimally different *Bill was sick last weekend* along with two analogous potential alternatives, *Bill was sick last Saturday* and *Bill was sick last Sunday*. In this case, too, it seems impossible to choose between the alternatives in a non-arbitrary way, but here (11) is of little help: while it may be impossible for one person to die on two different days, it is certainly possible for a person to be sick on two different days. A more general notion of symmetry should be able to capture those cases in which overlap does not help in making a non-arbitrary choice.

Finally, (11) considers only the case in which both alternatives are stronger than S. Consider the sentence *John died peacefully last weekend*, another minimal variant on our earlier example. We may again wish to consider the two potential alternatives *John died last Saturday* and *John died last Sunday*, and we would have no more grounds for choosing one alternative over the other in this case than in the context of the original *John died last weekend*. But differently from the original setting, (11) does not cover the present case: each of the potential alternatives allows for the possibility that John died on the relevant day but not peacefully, so neither alternative is stronger than the current assertion. A third natural extension of (11), then, would capture not just cases in which the alternatives are all stronger than S but also cases in which some of them are logically independent of it.

We can of course combine any two of the above extensions of the notion of symmetry, or indeed all three of them at once, and the problem of non-arbitrary choices will remain without change. Let us turn to an example that shows that the three extensions of (11) that we have just considered are indeed relevant, and that only F can eliminate dilemmas of this generalized kind. Suppose that artists are all musicians, painters, or

sculptors, and that it is possible to be an artist in more than one domain. Suppose further that there are tall and non-tall artists in every domain or combination of domains. We may now consider a sentence such as (31) along with three potential alternatives listed in (32).

(31) John met a tall artist

(32) Potential alternatives:
   a. John met a musician
   b. John met a painter
   c. John met a sculptor

The sentence in (31) and its potential alternatives in (32) illustrate all three generalizations of (11) at once: there are three potential alternatives, rather than two; there is overlap between the alternatives; and the alternatives are logically independent of the assertion. As just discussed, despite the deviation from the terms of (11), there seems to be no non-arbitrary way to choose between the alternatives, while trying to negate all of them together will result in contradicting the assertion. As the following examples show, in cases like this we obtain neither SIs nor *only*-inferences based on the alternatives under consideration:

(33) John met a tall artist
   *SI: ¬ John met a musician
   *SI: ¬ John met a painter
   *SI: ¬ John met a sculptor

(34) John only met a tall artist
   $\not\Rightarrow$ ¬ John met a musician
   $\not\Rightarrow$ ¬ John met a painter
   $\not\Rightarrow$ ¬ John met a sculptor

As before, we can use embedding under a universal operator to eliminate the contradiction and convince ourselves that *successful artist* indeed has *musician*, *painter*, and *sculptor* as alternatives:

(35) Every person in this room has met a tall artist
   SI: ¬ Every person in this room has met a musician
   SI: ¬ Every person in this room has met a painter
   SI: ¬ Every person in this room has met a sculptor

(36) John is required to marry a tall artist
   ⇒ ¬ John is required to marry a musician
   ⇒ ¬ John is required to marry a painter
   ⇒ ¬ John is required to marry a sculptor

It seems, then, that when generalized symmetry arises, no SIs or *only*-inferences using the participating alternatives are obtained. As in the narrow definition of symmetry in (11), the generalized notion of symmetry that seems to play a role in (33) and (34) involves alternatives that cannot be negated together without contradicting the

prejacent and for which any choice among the alternatives would seem arbitrary; differently from the narrow definition, however, the generalized notion allows for more than two alternatives, for alternatives that are not mutually exclusive, and for alternatives that are logically independent of the prejacent. We will shortly see that context is incapable of breaking generalized symmetry. Before that, let us ignore contextual restriction for a brief while longer and try to sharpen our understanding of generalized symmetry by looking at a suggestive similarity between this notion and another notion that was developed in the literature on exhaustivity.

From the perspective of a line of research going back at least to Groenendijk and Stokhof (1984) and including, more recently, van Rooij and Schulz (2004), Sauerland (2004b), Sevi (2005), Fox (2007), and Spector (2007), the absence of contradictory inferences in (33) and (34) – despite the potential for such inferences provided by the alternatives – is not surprising. From this perspective, exhaustivity inferences of the kind involved in SIs and in AF avoid at least some contradictions. In particular, Fox (2007) argues for a notion of exhaustivity that not only attempts to avoid contradictions but also avoids making arbitrary choices among the alternatives. Fox's definition involves the following notion, which he calls *innocent exclusion*:[14]

(37) Innocent exclusion: q is *innocently excludable* given p and a set A if q is in every maximal subset of A that can be negated consistently with p, $q \in IE(p, A)$

a. $IE(p, A) := \bigcap\{B \subseteq A : B \text{ is a maximal set in } A \text{ s.t. } \neg B \cup \{p\} \text{ is consistent }\}$

b. $\neg B := \{\neg p : p \in B\}$

An alternative is innocently excludable if it does not participate in dilemmas of the kind we have been considering, where the negation of all the alternatives will contradict the prejacent and a partial selection among the alternatives would be arbitrary. In (33) and (34) above, for example, none of the alternatives in (32) were innocently excludable. In (35) and (36), all of the embedded versions were. We can now state the general form of symmetry using innocent exclusion:

(38) Generalized symmetry: A set A of sentences is a generalized symmetric set of alternatives with respect to S if no element of A is in $IE(S, A)$

If innocent exclusion is indeed at the heart of SI and AF, the empirical pattern we just saw would make sense. Except, that is, for the question of contextual restriction. If context could break generalized symmetry, the following would be surprising:

(39) John met a (tall) musician; Sue met a (tall) painter; and Kim met a tall artist
*SI: ¬ Kim met a musician
*SI: ¬ Kim met a painter

Whatever (39) might mean, it does not license the inferences that Kim met neither a musician nor a painter, which, given our assumptions, would entail that Kim met a

[14] As Gajewski (2009a) notes, innocent exclusion does not always avoid contradiction. This is fortunate given the arguments in Fox and Hackl (2006) that exhaustification sometimes does give rise to contradiction, and Gajewski shows that in those cases innocent exclusion correctly predicts the relevant contradiction. In the cases discussed in the present paper, contradiction does not arise, and innocent exclusion derives this fact correctly.

sculptor. Again, embedding under a universal operator shows that the problem is not with an absence of alternatives:

(40) John is required to marry a (tall) musician; Sue is required to marry a (tall) painter; and Kim is required to marry a tall artist
SI: ¬ Kim is required to marry a musician
SI: ¬ Kim is required to marry a painter

We can conclude that, both in SI and in AF, contextual restriction is incapable of breaking generalized symmetry. This is a somewhat surprising fact, and it suggests that it is not only that F is the same in both SI and AF but also that contextual restriction operates in the same way in both cases. This, in turn, brings us closer to a unification of the two processes.

What might explain the inability of context to break generalized symmetry? Fox and Katzir (2011) showed that the inability of context to break *narrow* symmetry, as defined in (11), can be derived from considerations of relevance. On assumptions made elsewhere in the literature, relevance is closed under Boolean operations.[15] Fox and Katzir note that, if contextual restriction intersects F with the set of *relevant* alternatives (rather than just any set of alternatives, such as the often smaller set of salient alternatives), the assumption of closure under Boolean operations (along with the assumption that the prejacent is relevant) derives the inability of contextual restriction to break narrow symmetry. As also noted there, however, deriving similar results for some of the generalizations of symmetry using considerations of relevance is quite challenging. I will not attempt to address this challenge here. Instead, I would like to discuss a different perspective, suggested to Danny Fox and me by Emmanuel Chemla and Benjamin Spector (pers. comm.). Chemla and Spector note that we already have a mechanism that is incapable of breaking symmetry: exhaustivity operations, on the assumption mentioned above that such processes negate only innocently excludable alternatives, never negate any elements that are part of a symmetric subset. Suppose, then, that we have the following order of operations when S is exhaustified: (a) $F(S)$ is defined as above; (b) for every innocently excludable alternative $A \in IE(S, F(S))$, $\neg A$ is generated as a potential inference; (c) context selects from within this family of potential inferences. On this architecture, any symmetric subset of F will be ignored by step (b) and will not lead to any potential inferences. Consequently, contextual restriction in step (c) will never have the opportunity to break symmetry. No special restrictions on C are needed.

If Chemla and Spector's proposal is correct, we can define SI and *only* as follows. For convenience, let us start by defining $EXC(S)$, the conjunction of what gets excluded given S in a context C. The elements that are excluded are those that C selects from among the innocently excludable formal alternatives of S. The SI of S is simply $EXC(S)$ in the given context. The exhaustive operator *only* affirms the prejacent S and conjoins it with $EXC(S)$.[16]

[15] See Groenendijk and Stokhof (1984) and Lewis (1988), as well as the discussion in von Fintel and Heim (1997) as reported in Fox (2007).

[16] As mentioned earlier, this ignores important issues of presupposition and assertion along with the question of the height at which *only* operates. Such questions are not directly relevant for the present discussion.

(41) $EXC(S) = \bigwedge\{\neg S_i : S_i \in IE(S, F(S)) \;\&\; S_i \in C\}$

a. $SI(S) = EXC(S)$

b. $Only(S) = S \wedge EXC(S)$

6 Contrast

In section 4 we saw that the systematic absence of certain inferences (specifically, inferences involving non-monotone substitutions) seemed to favor a semantic restriction of the alternatives over the relatively permissive Complexity, defined in (25). Upon closer examination, we saw that the problem was due not to an absence of alternatives but rather to a surfeit of alternatives and the role played by symmetry and contradiction. In this section we will encounter another case of systematically absent inferences (specifically, inferences involving non-contrasting substitutions), and we will again be faced with the choice between accounting for the pattern in terms of too few alternatives or too many alternatives. As with monotonicity, we will see that the problem is due to a surfeit of alternatives and the effects of (generalized) symmetry and contradiction. An added value, in the present case, is that the examination will help clarify the roles of symmetry and contradiction for FF, an area where we have so far seen no way to probe these matters.

6.1 Wagner's puzzle

Wagner (2005, 2006, 2012) notes that permissive local alternatives, as in Complexity , lead to puzzling results with respect both to the semantics of AF and to the pragmatics of FF. Consider the example of AF in (42), in which focus falls on the adjective *red*. While Complexity would allow adjectives such as *red*, *blue*, and *expensive* to serve as alternatives to one another, (42) entails that John does not own blue convertibles but says nothing about whether he owns expensive convertibles. That is, (43a) appears to be negated and thus seems to be a well-behaved alternative to the prejacent in (42), while (43b) seems to be ignored.[17]

(42) John only owns red$_F$ convertibles

(43) a. John owns blue convertibles

b. John owns expensive convertibles

Similarly, consider the example of FF in (44):

(44) Mary's uncle, who is very rich and makes expensive convertibles, came to Mary's wedding. I wonder what he brought as a present.

a. He brought a cheap$_F$ convertible

b. # He brought a red$_F$ convertible

[17]While contextual restriction will not help in distinguishing the two hypotheses compared here, it is of course still operative. To simplify the discussion below, all the examples should be understood as being uttered in a context where the alternatives discussed in the text are salient.

Given the context-setting sentence in (44), (44a) is appropriate while (44b) is not. As noted by Wagner and others (see in particular Spathas (2010)), (44a) and (44b) are both meant to deny an expectation that is accommodated with the help of the context-setting (44). In this case, the expectation is that Mary's uncle brought an expensive convertible. Call this expectation η. According to COMPLEXITY, (44a) and (44b) both have η as an alternative, which means that both should be acceptable according to the appropriateness condition (4). Only the former, however, behaves as expected. This suggests that while *cheap* has *expensive* as an alternative, *red* does not.

6.2 Contrasting alternatives

Based on (42) and (44), Wagner argues for a restrictive view on which true focus alternatives – that is, the alternatives that end up being relevant for AF and for the appropriateness condition for FF – must be *contrastive* in the context of their sister node. Two elements are contrastive, according to Wagner, if they denote distinct cells in a salient partition. In particular, they must be mutually exclusive.

(45) A node α' *contrasts* with a node α in the context of a sister node β only if $[\![\ [\alpha'\ \beta]\ ]\!] \Rightarrow \neg [\![\ [\alpha\ \beta]\ ]\!]$.[18]

(46) *Allowable substitutions* (CONTRAST) elements that contrast with the target

In (42) above, for example, *blue* is a contrasting alternative to *red* in the context of the sister *convertible* since *blue convertibles* excludes *red convertibles*. According to CONTRAST, then, *blue* is a true alternative to *red* in this context. On the other hand, *expensive convertibles* does not exclude *red convertibles*, so *expensive* is not considered by CONTRAST to be a true alternative to *red* in the context of the sister *convertibles*. Consequently, when the operator *only* negates the alternatives to the prejacent, (43a) is negated and (43b) is not. Similar remarks apply for the FF example in (44).

6.3 AF: deriving the basic pattern with COMPLEXITY

The assumption that true alternatives must be mutually exclusive raises the question of how to account for scalar alternatives such as *some* and *every* (or *or* and *and*), which are not mutually exclusive and yet serve as alternatives of one another: *John did some of the homework* has the SI that John did not do all of the homework, and *John only did some$_F$ of the homework* has this inference as an entailment. This is a puzzling state of affairs, and in the remainder of this section I will try to show that the puzzle is only illusory: I will argue that CONTRAST, as defined in (46), can and should be avoided and that COMPLEXITY, as defined in (25), should be maintained. I will proceed along the same lines as the argument against adding the monotonicity requirement that we saw in section 4. Contrast and monotonicity are independent – some pairs (*some* and *none*; *or* and *nor*) contrast but are not co-monotonic, while other pairs (*some* and *all*; *or* and *and*) are co-monotone but do not contrast – but both conditions deal with their respective empirical puzzles in terms of too few alternatives. COMPLEXITY, on the

[18] Where $\Rightarrow$ is cross-categorial implication (for types that end in t) and $\neg$ is cross-categorial negation (again, for types that end in t). See Katzir (2013) for further details.

other hand, accounts for both puzzles in terms of too many alternatives, relying on contradiction and exclusion to explain the absence of the relevant inferences. As in the case of monotonicity, we will see that we can make the predictions of contrast diverge from those of COMPLEXITY by eliminating contradiction. I will try to show that in those cases, the facts support COMPLEXITY.

The argument presented below, from Katzir (2013), will have the following benefit for the present discussion: in addition to further supporting COMPLEXITY for the case of AF, it will shift the account of FF from one that relies on the appropriateness condition in (4) to one that relies on exhaustification. This, in turn, will allow us to explore the role of symmetry and contradiction in FF, a question that has been more difficult to address than its counterpart for SI and AF. My tentative conclusion will be that in FF, too, F is defined by COMPLEXITY while contextual restriction cannot break symmetry.

Let us start with the case of AF, where the argument will directly follow the lines of our earlier argument against monotonicity. Consider again (42) above, repeated here:

(47) John only owns red$_F$ convertibles

Let us adopt Wagner's assumption that color adjectives, such as *red* and *blue*, induce one partition of convertibles and that price adjectives, such as *cheap* and *expensive*, induce a different partition of convertibles that is independent of the first partition: a red convertible may be cheap, and it may be expensive, and the same holds for a blue convertible. To simplify the discussion, let us further assume that the set of adjectives is limited to {*red*, *blue*, *cheap*, *expensive*}.[19] This means that, on COMPLEXITY, the set of alternatives to which the operator *only* has access are the following:

(48) a. John owns blue convertibles
b. John owns red convertibles
c. John owns cheap convertibles
d. John owns expensive convertibles

In the current case, given the prejacent (48b) and the alternatives in (48), negating (48c) entails that (48d) is true: if John owns red convertibles and does not own cheap convertibles, then he owns expensive convertibles. Similarly, negating (48d) entails that (48c) is true. The choice between negating (48c) and negating (48d) is thus arbitrary. Consequently, neither alternative is innocently excludable in the sense of (37); assuming that *only* is defined as in (41b), neither alternative will be negated. On the other hand, negating (48a) leads to no arbitrary conclusions: if John owns red convertibles and does not own blue convertibles the truth of the remaining alternatives remains undetermined (we still do not know whether he owns cheap convertibles, and we do not know whether he owns expensive ones). Consequently, (48a) is innocently excludable and will be negated. By taking innocent exclusion into account, then, we derive the correct inferences for (47) using COMPLEXITY.

[19] The assumptions regarding the possible adjectives and their properties simplify in several different ways. For example, there are many additional color adjectives, and there may be additional price ranges to consider. We may also consider it possible for a convertible to have more than one color or to have no color. I hope that the present discussion can be extended to more realistic assumptions.

6.4 AF: distinguishing COMPLEXITY and CONTRAST

What allowed COMPLEXITY to derive the same predictions as CONTRAST in (47) is the fact that it is contradictory to own red convertibles and yet not own cheap convertibles and not own expensive convertibles. If we could modify (47) so as to eliminate the contradiction between the prejacent and the negation of the two alternatives with 'cheap convertibles' and 'expensive convertibles', the predictions of the two approaches will diverge. COMPLEXITY predicts that those alternatives will now be negatable. CONTRAST, on the other hand, predicts that the modified version will exhibit the same pattern of negation: since neither *cheap convertible* nor *expensive convertible* excludes *red convertible*, neither will be a true alternative to it.

One way to eliminate the contradiction is to change the verb in the relevant examples. The contradiction allowing COMPLEXITY to account for the entailments of (47) depends crucially on inferences that stem from lexical properties of the verb *own*: if one doesn't own a certain kind of convertible, then one doesn't own any instance of that kind of convertible. This property (related to the extensionality of *own*) is not shared by all transitive verbs. The intensional verb *collect*, for example, exhibits a different pattern of entailment (other verbs that pattern with *collect* include *seek*, *wish for*, and *crave*): it is possible to collect red convertibles and yet to not collect cheap convertibles and not collect expensive convertibles. If we replace *own* with *collect* in (47), then, the contradiction will disappear. COMPLEXITY predicts that the potential alternatives with *cheap* and with *expensive* will now be negated. CONTRAST, on the other hand, predicts as before that these will not be true alternatives. The facts support COMPLEXITY: (49) entails that John does not collect blue convertibles, that he does not collect cheap convertibles, and that he does not collect expensive convertibles; that is, it entails that (50a), (50c), and (50d) are all false.

(49) John only collects red_F convertibles

(50) a. John collects blue convertibles
b. John collects red convertibles
c. John collects cheap convertibles
d. John collects expensive convertibles

A different way to make the same point is to keep the original verb and embed it under a universal operator such as *require*. While it is contradictory to own red convertibles and yet not own cheap convertibles and not own expensive convertibles, there is nothing contradictory about being required to own red convertibles (say, in order to be admitted into a club) and yet not being required to own cheap convertibles and not being required to own expensive convertibles. COMPLEXITY, then, predicts that (51) will entail the negation of (52a), (52c), and of (52d), which it does. As before, CONTRAST makes the incorrect prediction that only the contrasting (52a) is a true alternative.[20]

(51) John is only required to own red_F convertibles

[20] At this point it might seem tempting to change the domain of locality at which contrast is determined. See Katzir (2013) for an argument that this will not help and that it is the contrast requirement itself that should be removed.

(52) a. John is required to own blue convertibles
b. John is required to own red convertibles
c. John is required to own cheap convertibles
d. John is required to own expensive convertibles

6.5 Free focus

We just saw that innocent exclusion allows Complexity to account for Wagner's puzzle in the case of AF. Moreover, we saw that when contradiction is eliminated, the predictions of Complexity and Contrast diverge and those of Complexity are borne out. In the present section I will try to show that a similar argument can be made in the case of FF.

Consider again (44) above, repeated here:

(53) Mary's uncle, who is very rich and makes expensive convertibles, came to Mary's wedding. I wonder what he brought as a present.
a. He brought a cheap$_F$ convertible
b. # He brought a red$_F$ convertible

As discussed in section 6.1, Complexity considers the accommodated expectation η = *that Mary's uncle brought an expensive convertible* to be an alternative both of (53a) and of (53b), which makes the contrast between the acceptability of (53a) and the oddness of (53b) look surprising. For Contrast, on the other hand, η is a contrasting, and hence true alternative of (53a) but not of (53b), thus predicting the contrast on the basis of the appropriateness condition in (4).

(53) has no occurrence of the operator *only*, but I would like to suggest that the explanation for the contrast in acceptability between (53a) and (53b) involves exhaustification and innocent exclusion nonetheless. Let us look first at the acceptable (53a). Without exhaustification, this response is irrelevant to the evaluation of η (it is possible to bring two convertibles, a cheap one and an expensive one, at the same time). If (53a) is exhaustified, on the other hand, we obtain the entailment that Mary's uncle did not bring an expensive convertible. This entailment addresses η, which in turn makes it an appropriate response in the given context.

Turning now to (53b), we can again see that without exhaustification the utterance is irrelevant to the evaluation of η. In this case, however, adding exhaustification is of little help: both (54c) and (54d) are alternatives, and since negating both contradicts the prejacent (54b) and negating just one will be arbitrary, neither of them is innocently excludable; consequently, assuming that exhaustification operates as defined in (41), neither will be negated. The result of exhaustification, then, is as irrelevant to η as the original (53b).

(54) a. He brought a blue convertible
b. He brought a red convertible
c. He brought a cheap convertible
d. He brought an expensive convertible

Note that the requirement that exhaustification be prohibited from breaking generalized symmetry – a requirement that follows from definition (41) above – is important. If context were not limited in this way, we would expect that when (53b) was uttered, context would be able to keep (54d), which was mentioned in the context provided in (53), and leave out (54c), which was not mentioned. This, in turn, would have made the exhaustification of the prejacent entail that Mary's uncle did not bring an expensive convertible, which is again relevant for addressing η, thus eliminating our explanation for the contrast between (53a) and (53b). The current account, then, relies on a notion of exhaustification that is very similar to that argued for in the case of SI and AF. In particular, this exhaustification uses innocent exclusion and does not allow context to break generalized symmetry.

If contradiction and exhaustification are indeed the source of the unacceptability of (53b), we again derive the prediction that eliminating contradiction should allow the non-contrastive alternatives to emerge. And again, Contrast predicts that such alternatives will not be allowed. As before, I will use both the verb *collect*, which does not share the entailments of *bring*, and embedding under the universal operator *require* to test the predictions of the two theories.

(55) The people in this club are very particular about the cars they collect. Mary, for example, collects expensive convertibles.
   a. And John collects cheap$_F$ convertibles
   b. And John collects red$_F$ convertibles

(56) Mary was required to bring an expensive convertible.
   a. And John is required to bring a cheap$_F$ convertible
   b. And John is required to bring a red$_F$ convertible

(55a) and (55b) are both acceptable responses to (55), and (56a) and (56b) are both acceptable responses to (56). We can conclude that in FF, too, eliminating contradiction allows non-contrastive alternatives to be negated. Again, this is as expected by Complexity but it is surprising under Contrast.[21]

We have just seen a defense of Complexity against Contrast, a third semantic account (following the discussion of Types and Monotonicity earlier in the paper). More interesting from the present perspective is that in addressing Wagner's puzzle we now have a way to discuss the roles of symmetry and contradiction for FF. Recall that Wagner's account for the case of FF relies on the appropriateness condition in (4). If we could maintain an account in terms of the appropriateness condition, there would be no reason to think that considerations of symmetry apply to FF. In particular, one could imagine that an account of FF in terms of semantic type could be maintained. What we arrived at, however, is an account that relies on a consideration that is quite different from the appropriateness condition, namely the ability of an utterance to address an expectation. Crucially, this required that we take into account the exhaustified meaning of utterances. In previous sections we examined exhaustification in SI and AF and saw reasons for thinking that the breaking of symmetry is performed in F, using a

[21] As in the case of AF, the problem is with the contrast requirement itself and not with the height at which it is applied. See Katzir (2013) for discussion.

structural definition of allowable substitutions, and that context is limited to choosing from innocently excludable alternatives. By analyzing FF as involving exhaustification in which context can only choose from among innocently excludable alternatives, we now have reason to take the division of labor in FF between F and C to be the same with respect to symmetry-breaking as in SI and AF, which suggests a unified treatment of all three processes.

7 Summary

I have sketched a unified view of SI, AF, and FF that has the following properties:

(57) a. In the division between formal alternatives and contextual restriction, only the former can break symmetry

b. The formal alternatives of an object are defined through markedness-based substitutions within focused constituents, where markedness is cashed out in terms of structural complexity

c. All three phenomena rely on exhaustification in which context selects from within those formal alternatives that are innocently excludable given the assertion

To argue for (57a), I started by reviewing the reasons to think that symmetry is a problem for the type-based approach in the case of SI and AF: this approach leaves all symmetry breaking to context, and we saw that this choice is problematic. We later saw that, both for SI and for AF, only F can ever break symmetry: when we forced F to include a symmetric subset, this symmetry could not be broken by other processes.

To argue for (57b), we saw that, once entailment relations are taken into account, strictly simpler alternatives emerge quite generally. We also saw that alternatives are asymmetric: if ϕ is strictly simpler than ψ, then usually ϕ will be an alternative of ψ but ψ will not be an alternative of ϕ. We then saw that the systematic absence of certain inferences seemed to motivate two further restrictions of the allowable substitutions: monotonicity and contrast. Upon closer inspection, though, we saw that the unrestricted structural account had its own way of accounting for the missing inferences, using contradiction and symmetry. By eliminating the contradiction we allowed the predictions of the two approaches to diverge, and those of the unrestricted structural account were borne out by the data.

Finally, FF required shifting from the appropriateness condition to an account that relies on exhaustification, supporting (57c). Since exhaustification in FF exhibited similar properties with respect to the breaking of generalized symmetry as in SI and AF, we obtained evidence for unifying FF with SI and AF.

References

Anderson, Stephen. 1969. West Scandinavian vowel systems and the ordering of phonological rules. Doctoral Dissertation, MIT, Cambridge, Mass.

Atlas, Jay David, and Stephen Levinson. 1981. It-clefts, informativeness, and logical form. In *Radical pragmatics*, ed. Peter Cole. Academic Press New-York.

Büring, Daniel. 2012. What's given (and what's new) in the theory of focus. In *Proceedings of the Thirty-Fourth Annual Meeting of the Berkeley Linguistics Society, February 8–10 2008*, ed. Sarah *et al.* Berson, 403–424. Berkeley, CA: Berkeley Linguistics Society.

Chomsky, Noam. 1995. *The minimalist program*. Cambridge, MA: MIT Press.

von Fintel, Kai, and Irene Heim. 1997. Classnotes on pragmatics. MIT.

Fox, Danny. 2000. *Economy and semantic interpretation*. Cambridge, MA: MIT Press.

Fox, Danny. 2007. Free choice disjunction and the theory of scalar implicatures. In *Presupposition and implicature in compositional semantics*, ed. Uli Sauerland and Penka Stateva, 71–120. Palgrave-Macmillan.

Fox, Danny, and Martin Hackl. 2006. The universal density of measurement. *Linguistics and Philosophy* 29:537–586.

Fox, Danny, and Roni Katzir. 2011. On the characterization of alternatives. *Natural Language Semantics* 19:87–107.

Gajewski, Jon. 2009a. Innocent exclusion is not contradiction free. Ms., UConn.

Gajewski, Jon. 2009b. L-triviality and grammar. Handout, UConn Logic Group, 2/27/09, February 2009.

Gazdar, Gerald. 1979. *Pragmatics: Implicature, presupposition and logical form*. Academic Press New-York.

Geurts, Bart. 2010. *Quantity implicatures*. Cambridge University Press.

Grice, Paul. 1989. *Studies in the way of words*. Cambridge, MA: Harvard University Press.

Groenendijk, Jeroen, and Martin Stokhof. 1984. Studies in the semantics of questions and the pragmatics of answers. Doctoral Dissertation, Universiteit van Amsterdam, Amsterdam.

Halle, Morris. 1997. Distributed Morphology: Impoverishment and fission. In *Papers at the interface*, ed. B. Bruening, Y. Kang, and M. McGinnis, 125–149. Cambridge, Mass.: MIT Working Papers in Linguistics.

Halle, Morris, and Alec Marantz. 1993. Distributed morphology and the pieces of inflection. In *The view from Building 20*, ed. Kenneth Hale and Jay Keyser, 111–176. Cambridge, MA: MIT Press.

Hamblin, C. L. 1973. Questions in Montague English. *Foundations of Language* 10:41–53.

Hirschberg, Julia. 1985/1991. *A theory of scalar implicature*. New York: Garland. PhD Thesis, University of Pennsylvania.

Horn, Laurence. 1972. On the semantic properties of the logical operators in English. Doctoral Dissertation, UCLA.

Horn, Laurence. 1978. Lexical incorporation, implicature, and the least effort hypothesis. In *Papers from the Parasession on the Lexicon*, 196–209. Chicago: Chicago Linguistics Society.

Horn, Laurence. 1984. Toward a new taxonomy for pragmatic inference: Q-based and R-based implicatures. In *Meaning, form, and use in context*, ed. D. Schiffrin, 11–42. Washington: Georgetown University Press.

Horn, Laurence. 1989. *A natural history of negation*. Chicago: University of Chicago

Press.
Horn, Laurence. 2000. From IF to IFF: Conditional perfection as pragmatic strengthening. *Journal of Pragmatics* 32:289–326.
Jackendoff, Ray. 1972. *Semantic interpretation in generative grammar*. Cambridge, MA: MIT Press.
Katzir, Roni. 2007. Structurally-defined alternatives. *Linguistics and Philosophy* 30:669–690.
Katzir, Roni. 2013. A note on contrast. *Natural Language Semantics* (In press).
Katzir, Roni, and Raj Singh. 2013. Constraints on the lexicalization of logical operators. *Linguistics and Philosophy* 36:1–29.
Kiparsky, Paul. 1973. "elsewhere" in phonology. In *A festschrift for Morris Halle*, ed. Stephen Anderson and Paul Kiparsky, 93–106. Holt, Rinehart and Winston, Inc.
Kroch, Anthony. 1972. Lexical and inferred meanings for some time adverbials. Quarterly Progress Reports of the Research Laboratory of Electronics 104, MIT, Cambridge, Mass.
Lewis, David. 1988. Relevant implication. *Theoria* 54:161–174.
Matsumoto, Yo. 1995. The conversational condition on Horn Scales. *Linguistics and Philosophy* 18:21–60.
McCawley, James D. 1978. Conversational implicatures and the lexicon. In *Syntax and semantics*, ed. Peter Cole, volume 9: Pragmatics, 245–259. New York: Academic Press.
Poser, William J. 1992. Blocking of phrasal constructions by lexical items. In *Lexical matters*, ed. Ivan Sag and Anna Szabolcsi, 111–130. Stanford, CA: CSLI Publications.
van Rooij, Robert, and Katrin Schulz. 2004. Exhaustive interpretation of complex sentences. *Journal of Logic, Language and Information* 13:491–519.
Rooth, Mats. 1985. Association with focus. Doctoral Dissertation, University of Massachusetts, Amherst.
Rooth, Mats. 1992. A theory of focus interpretation. *Natural Language Semantics* 1:75–116.
Sauerland, Uli. 2004a. On embedded implicatures. *Journal of Cognitive Science* 5:107–137.
Sauerland, Uli. 2004b. Scalar implicatures in complex sentences. *Linguistics and Philosophy* 27:367–391.
Sauerland, Uli. 2005. Don't interpret focus! Why a presuppositional account of focus fails, and how a presuppositional account of givenness works. In *Proceedings of Sinn und Bedeutung 9*, ed. Emar Maier, Corien Bary, and Janneke Huitink, 370–384. Nijmegen: University of Nijmegen.
Schwarzschild, Roger. 1999. GIVENness, AvoidF and other constraints on the placement of accent. *Natural Language Semantics* 7:141–177.
Sevi, Aldo. 2005. Exhaustivity: A semantic account of 'quantity' implicatures. Doctoral Dissertation, Tel-Aviv University.
Spathas, Giorgos. 2010. Focus on anaphora: Accent placement and the syntax and semantics of anaphors. Doctoral Dissertation, Universiteit Utrecht.
Spector, Benjamin. 2007. Scalar implicatures: Exhaustivity and Gricean reasoning. In *Questions in dynamic semantics*, ed. Maria Aloni, Alistair Butler, and Paul Dekker,

229–254. Elsevier.
Swanson, Eric. 2010. Structurally defined alternatives and lexicalizations of XOR. *Linguistics and Philosophy* 33:31–36.
Wagner, Michael. 2005. Prosody and recursion. Doctoral Dissertation, MIT, Cambridge, Mass.
Wagner, Michael. 2006. Givenness and locality. In *Proceedings of SALT 16*, ed. Jonathan Howell and Masayuki Gibson, 295–312. Ithaca, NY: CLC Publications.
Wagner, Michael. 2012. Focus and givenness: a unified approach. In *Contrasts and positions in information structure*, ed. Ivona Kučerová and Ad Neeleman, 102–147. Cambridge: Cambridge University Press.
Zipf, George Kingsley. 1949. *Human behavior and the principle of least effort*. Cambridge, MA: Addison-Wesley.